UCRL-92792
PREPRINT

THE ROLE OF SUPERCOMPUTERS IN MAGNETIC FUSION AND ENERGY RESEARCH PROGRAMS

John Killeen

This paper was submitted for the Subcommittee on Energy Development and Applications and the Subcommittee on Science, Research & Technology, U.S. House of Representatives, Tallahasse, Florida, June 10, 1985.

June 1985

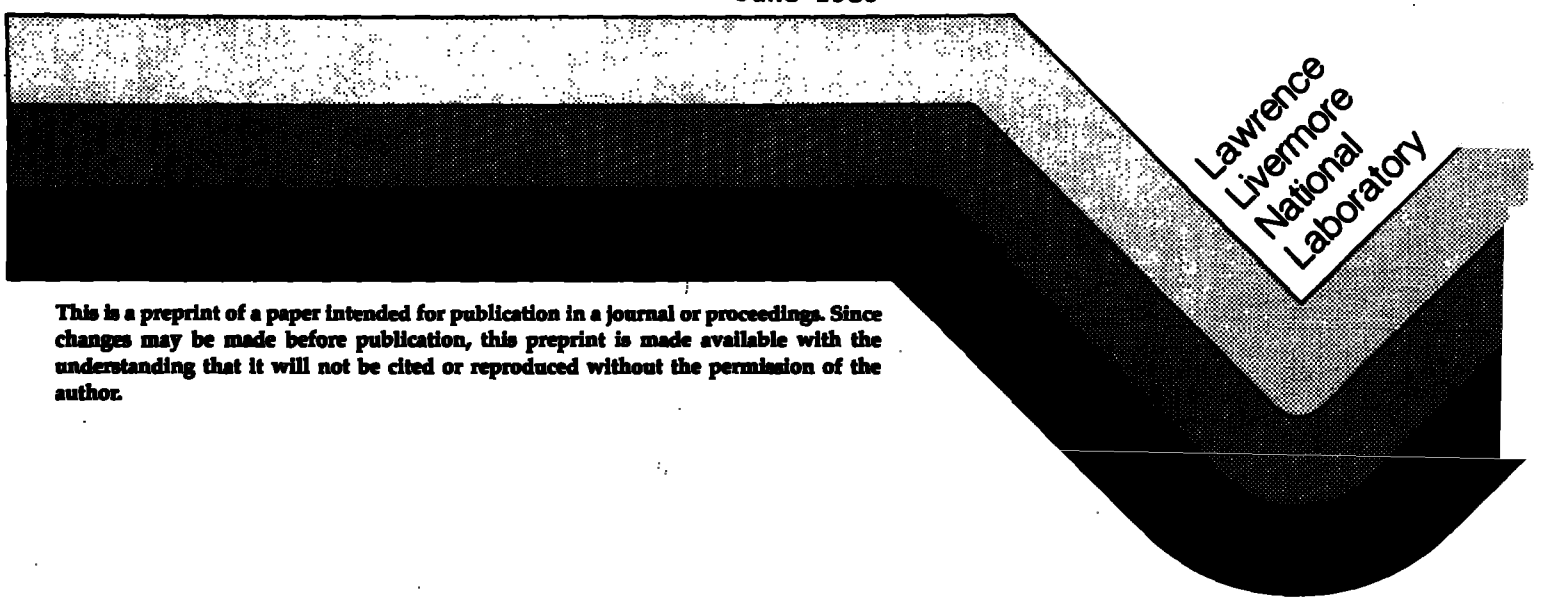

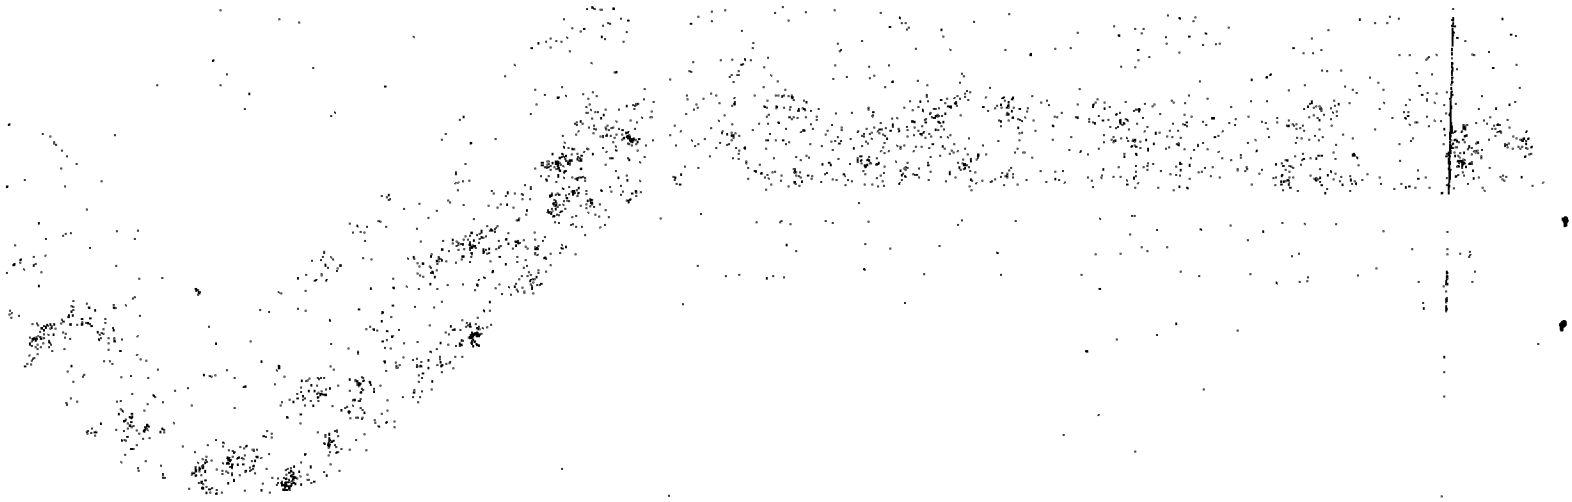

THE ROLE OF SUPERCOMPUTERS IN MAGNETIC FUSION AND ENERGY RESEARCH PROGRAMS *

John Killeen
National Magnetic Fusion Energy Computer Center
Lawrence Livermore National Laboratory
Livermore, CA 94550

Abstract

The importance of computer modeling in magnetic fusion (MFE) and energy research (ER) programs is discussed. The need for the most advanced supercomputers is described, and the role of the National Magnetic Fusion Energy Computer Center in meeting these needs is explained.

MAGNETIC FUSION

INTRODUCTION

During the early 1970's the U.S. magnetic fusion program supported at least fifteen varieties of experimental concepts. These were rather small experiments as compared to today's large facilities. During the years 1974 to 1980, the program went through a period of dramatic growth, but at the same time evaluations and reviews reduced the number of experimental concepts supported to the following six:

Tokamak
Tandem Mirror
Reverse Field Pinch
Stellarator
Compact Toroids
Elmo Bumpy Torus

The most advanced of the above concepts is the Tokamak, and all four of the major international groups have commissioned large facilities to establish the scientific feasibility of fusion.

All of the international groups are designing forms of "The Next Step," which is an ignition Tokamak. A common feature of these designs is their large projected cost, so at the present time none have been authorized.

*Work performed under the auspices of the U.S. Department of Energy by the Lawrence Livermore National Laboratory under contract number W-7405-ENG-48.

A large tandem mirror experiment (MFTF-B) is being built in the U.S. Advanced stellarators are planned for Kyoto, Japan and Garching, FRG, and a large RFP is planned for Padua, Italy. New stellarator, RFP, and compact toroid facilities are being proposed in the U.S.

In all of these concepts, there are eight fusion physics issues which must be addressed as a complete plasma system, i.e., they are interdependent. They are:

MHD Equilibrium and Stability
Perpendicular Ion and Electron Confinement
Parallel Confinement
Electric Potential
Heating
Fueling
Impurity Influx
Alpha Particle Heating

In order to resolve these issues, i.e. to reach a state where a fusion reactor is feasible, the experimental programs must be augmented by a program of computer simulation to aid in the design and interpretation of the experiments and implementation of theory. The following plasma physics models are of importance to the fusion program.

Time-dependent magnetohydrodynamics
Plasma transport in a magnetic field
MHD and guiding-center equilibria
MHD stability of confinement systems
Vlasov and particle models
Multi-species Fokker-Planck codes
Hybrid codes

The need for such a variety of models is caused by the great variation in time and space scales[1] present in the plasma phenomena relevant to the eight fusion physics issues. The implementation of these models requires the most advanced

supercomputers available. The impact of new supercomputers on some of the types of models will be discussed later in this paper.

In addition to plasma physics models, advanced engineering computations must be made. Engineering models needed in fusion reactor design studies include:

Plasma engineering-burning plasma dynamics
Nucleonics
Mechanical design
Magnetic field analysis
Systems studies
Thermal hydraulics
Tritium handling
Safety and environmental studies

NEW SUPERCOMPUTERS

Supercomputers are the most powerful general-purpose computers available for information processing. Currently, supercomputers have the capability of performing hundreds of millions of arithmetic or floating point operations per second (MFLOPS) and are used in two general areas: real-time applications such as signal processing and scientific computing. In the race to build the next generation of supercomputers, scientists are experimenting with a variety of architectural designs. The new architectures will have as few as two processors with shared memories to extensive parallel architectures with hundreds of local memories and processors, all executing instructions simultaneously.

There are three types of parallel architecture capable of increasing performance by a hundredfold over today's state-of-the-art supercomputers. They are:

Lockstep vector processors,
Tightly coupled parallel processors,
Massively parallel devices.

When the execution unit operates simultaneously (in lockstep) on many data entities, the machine is said to have an array architecture. When the execution unit operates on sets of data, on an assembly line basis, the machine is termed a vector processor or pipeline processor. The CDC 205 and Cray I are examples of vector processors. The real beneficiaries of such vector processors have turned out to be multi-dimensional fluid codes, which are dominated by long vector loops.

A second architectural type capable of a hundredfold increase over state-of-the-art supercomputers is tightly coupled systems of a few high-performance processors. In principle, collaboration of these processors on a common task can produce the two orders of magnitude speedup that is needed.

The current trend in supercomputer architecture is toward tightly coupled systems with two to four vector processors typically sharing a large memory. Recent experiments suggest that these systems can be successfully used in parallel processing of scientific computations. The next logical step in this trend is toward systems with 8, 16, or more processors.

In the long term it is possible to build massively parallel systems, that is, systems with 1000 or more processors communicating with thousands of memories. In general, the scientist cannot manually find and manage parallelism for thousands of processors. Rather, the software must find it, map it onto the architecture, and manage it. Therein lies a formidable research issue for massively parallel computation.

The following two tables list (1) existing supercomputers, and (2) announced supercomputers. This tabulation employs only the few parameters usually contained in press-release-type information.

NATIONAL MFE COMPUTER CENTER

The MFE Computer Network (Figure 1) provides fusion researchers in the U.S. the full range of available computational power in the most efficient and cost effective manner. This is achieved by using a network of computers of

TABLE I Current Supercomputers

Organization	Fujitsu	Hitachi	CDC	CRAY	CRAY
Model	VP-200	S-810/20	205	X-MP/2	X-MP/4
announcement	Jul 1982	AUg 1982	Jun 1981	Aug 1982	Aug 1984
architecture (64 bit words)	vector (IBM compatible)	vector (IBM compatible)	vector	vector multi-processor 2 CPU	vector multi-processor 4 CPU
maximum performance (M FLOPS)	500	630	400	479	953
maximum main memory size (64 bit words)	32M MOS	32M MOS	16M MOS	4M Bipolar	8M Bipolar

TABLE II Supercomputers Now In Design

Organization	CRAY	CRAY	ETA	Denelcor	NEC
Model	2	3	GF10	HEP-2	SX-2
announcement (or project start)	1985	none officially	Sept 1983	May 1983	April 1983
availability	1985	1986	1986	1986	1985
architecture	vector multi-processor 4 CPU	vector multi-processor 16 CPU	vector multi-processor 8 CPU	scalar multi-processor 64 CPU	vector
maximum performance (M FLOPS)	1,000	10,000	10,000	4,000	1,300
maximum main memory size (64 bit words)	256M MOS	256M MOS	256M MOS	256M MOS	32M MOS

different capability tied together and to the users via dedicated data lines and dial up telephone lines. The concept of the NMFECC is that different levels of computer capability are provided at the various locations according to research priorities. At the national center (Figure 2), providing high level capability to the entire community, are two high-speed Cray 1 computers, and a Cray X-MP/2. Additional equipment at the national center includes processors and other ADP equipment for communications, file management, and data storage.

On May 28, 1985 the first Cray 2 computer system was delivered to the NMFECC. This computer has four vector processors and 64 million words of MOS memory. This system will give the fusion community the capability required for advanced plasma modeling as described in the next section.

At the next level of capability are User Service Centers (USC's): DEC-10 computer systems with direct high-speed access to the national center through PDP-11/40 remote communications control processors. There are now five operational USCs (Figure 1) in the field located at Princeton Plasma Physics Laboratory (PPPL), the Los Alamos National Laboratory (LANL), the Oak Ridge National Laboratory (ORNL), GA Technologies, Inc. (GA), and LLNL (for the mirror confinement program). A sixth USC, used in center operations, is located at the NMFECC itself.

A third level of capability is provided through the Network Access Port (NAP). MFECC designed the NAP to permit remote computers to be connected to the MFE network as remote hosts.

A fourth level of capability is provided by Remote User Service Stations (RUSS) at selected sites (Figure 1). RUSS stations provide users with the capability of printing output files locally on a 1000 line/minute printer and act as a terminal concentrator for up to 16 interactive terminal users. RUSS stations are connected to the nearest MFE-NETWORK communications processor over a 9600 baud dedicated line (Figure 1).

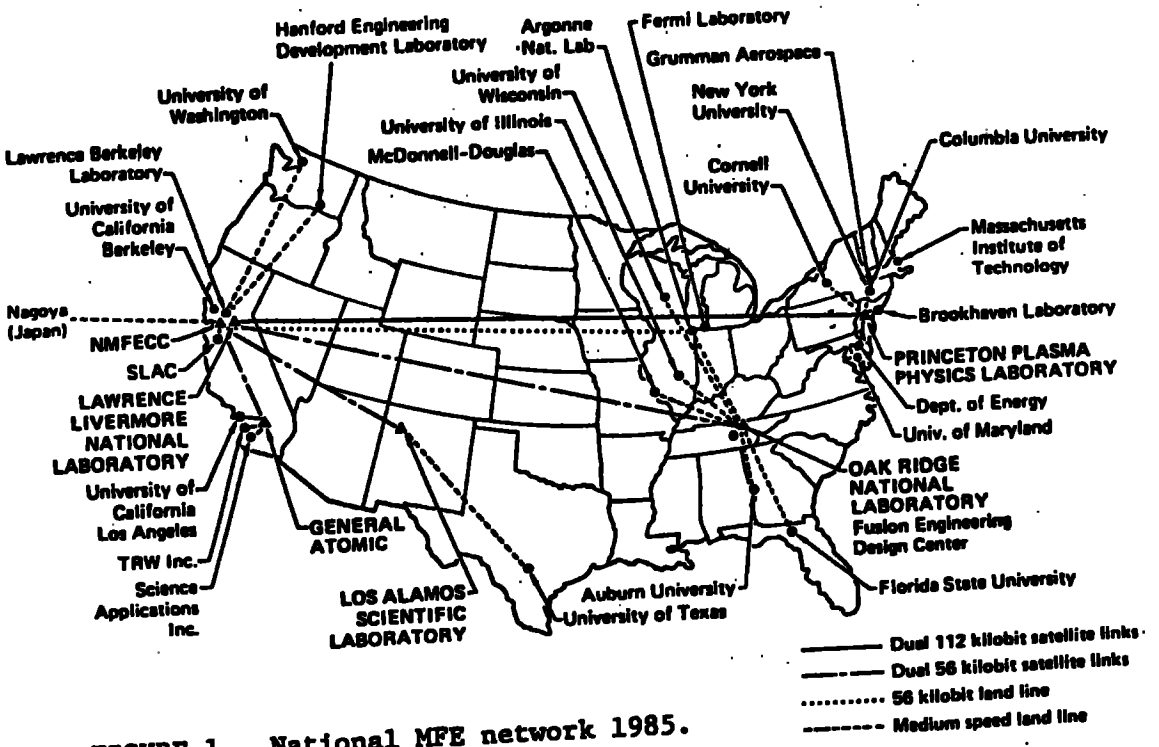

FIGURE 1 National MFE network 1985.

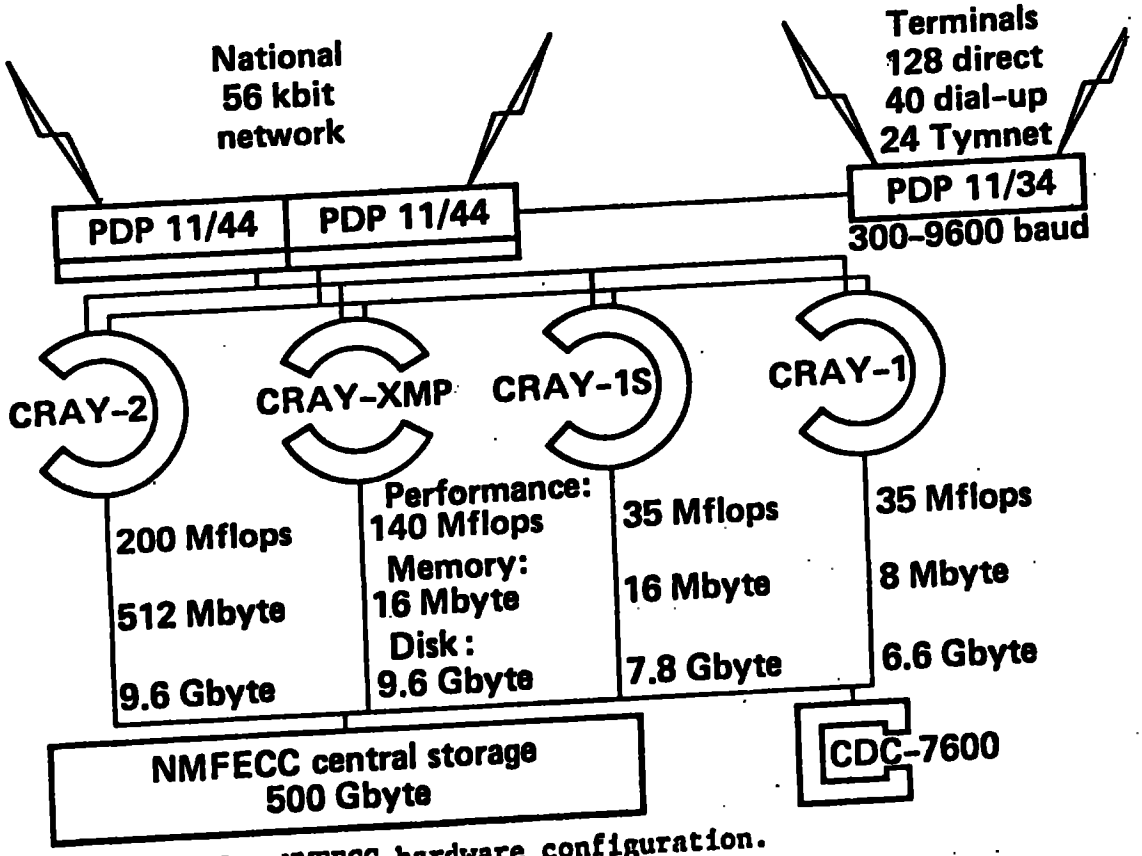

FIGURE 2 NMFECC hardware configuration.

Data-Communications Systems

Data Communications services to the National MFE Computer Center are provided on a 24 hours/7 day basis. Three types of service are provided to NMFECC users as outlined below:

1. Wide band Satellite Network Service. Users at Major USC's on the MFE net may log on to their local DEC-10 system and interact with the computing resources at the Central facility in Livermore.

2. Dedicated 9600 Baud Service. Remote User Service Stations on the MFE Net are served by dedicated leased 9600 baud lines which terminate either at the Center (LLNL) or at the nearest MFE Communications Control Processor (Figure 1).

3. Dial Up Service. Users not at major fusion laboratories may dial-up the Center using one of the following services: (a) TYMNET, (b) ARPANET and (c) DIRECT DIAL COMMERCIAL.

NMFECC Computing Environment

The NMFECC computing environment reflects the needs of computer users in the Magnetic Fusion Energy research community. Both interactive timesharing and batch processing are available. The fusion community has always found that interactive computing, even with the largest codes, is by far the most efficient use of physicists efforts. The 5% overhead in swapping codes in and out of the machines provides fast debugging, immediate turn around on key results, and the capability to interact with codes which need user control. The Livermore Time Sharing System (LTSS) was adapted by the NMFECC for the Cray 1 computer in about six months. CTSS is supported by libraries of FORTRAN callable subroutines which enable a user to issue almost every system call, giving access to every part of the hardware. A typical physics code can be run from a terminal, display graphics as it runs, be interrupted or interrogated at any time. The ability to start or stop a code at any point and inspect the results provides debugging at least 100 times faster than older methods. The CTSS operating system is also used on the Cray X-MP/2 and the Cray 2.

PLASMA MODELING IN MAGNETIC FUSION

It is within our grasp at present to model plasmas in full 3D with one or two orders of magnitude variation in space and time scales in each problem. Some of the recent success in the field are worth listing as they are the basis for further developments. Representative work in MHD, Kinetic models, and Fokker-Planck calculations are considered.

Time-dependent MHD Codes

A technique for determining MHD instabilities along with their growth rates is through the solution of the time dependent MHD equations of motion. The full set of MHD calculations comprise a coupled system of eight nonlinear partial differential equations, the solution of which is a formidable task on any computer system. In order to make these computations tractable, approximations have often been made, including reduction in dimensionality, linearization, restriction to a particular geometry, ordering, or regime, and the assumption of no transport or resisitivity.

The recent advances in three-dimensional resistive MHD calculations for tokamaks have depended crucially on obtaining a reduced set of MHD equations by expanding in the inverse aspect ratio[2]. This is possible because of the strong and almost uniform toroidal magnetic field in tokamaks. Additionally, the computational speed of the codes based on the tokamak reduced equations is greatly enhanced by the assumption of incompressibility, which eliminates the compressional Alfvén wave. Because of the strong field in a tokamak, the fastest remaining mode evolves on a time-scale on the order of the major circumference divided by the Alfvén velocity. This time scale may be more than an order of magnitude longer than that of the compressional Alfvén wave. Since the field components in the Reversed Field Pinch (RFP) are all of the same order, and since these devices possess finite beta, there exists no universally small parameter in which to expand the basic equations. Instead, the full equations are integrated, using care to separate compressible and incompressible motions as much as possible[3],[4]. These simulations reproduce some features of present experiments, but the next generation of computers is clearly needed here.

To make these three-dimensional codes applicable to more general geometries (e.g. stellarators) and to simultaneously include enough effects to ensure a complete description of the important physics effects (e.g. parallel heat transport, compressibility, finite larmor radius effects, and smaller values of resistivity) requires a machine with about 10 times the CPU speed of the Cray 1 as well as a large memory, e.g. the Cray 2.

Particle and Hybrid Codes

In many cases fluid models are not adequate to describe plasma behavior, for it is necessary to consider microscopic effects, i.e., the effects of the way particles are distributed in velocity. Numerically this is most often accomplished through particle codes[5-7]. Fully nonlinear kinetic ion and electron simulations in 2-D Cartesian geometry have been carried out over the last decade. In the past, Cartesian geometry was not a major physics limitation even with the obvious cylindrical and toroidal nature of experiments, because these models necessarily dealt with length and time scales on the order of the electron gyroradius and plasma oscillation period for stability. Resolving such length and time scales meant that any realistic macroscopic dimension could be considered infinite. With the increase of grid resolution allowed by improved computers and methodology, the scope of particle simulations has grown to encompass nonlocal effects and more realistic geometries.

On the present computers, large scale particle simulations in 2-1/2D and 3D are mainly limited by the size of the maximum fast memory of the Cray 1 (of the order of 1 M words, or 2 M for the Cray 1S). Experimentally relevant physics problems in magnetic confinement have important three-dimensional aspects, such as in the multiple-helicity interaction of collisionless tearing modes and in the drift wave turbulence in sheared magnetic fields; the 64M word memory of the Cray 2 and its vector addressing will greatly enhance these simulations.

Particle-fluid hybrid models have become important in the last five years. A typical hybrid model represents the ion components as kinetic species and the electrons as a fluid in order to eliminate some or all fast electron

frequencies and short length scales. Recent progress with hybrid models is impressive but is still quite computationally expensive (typically taking roughly two to four times more Cray CPU time than does an MHD code of equal dimensionality).

Fokker-Planck Codes

In the simulation of magnetically confined plasmas where the ions are not Maxwellian and where a knowledge of the distribution functions is important, kinetic equations must be solved. At number densities and energies typical of mirror machines, end losses are due primarily to the scattering of charged particles into the loss cones in velocity space by classical Coulomb collisions. The kinetic equation describing this process is the Boltzmann equation with Fokker-Planck collision terms. The heating of and current generation in plasmas by energetic neutral beams and microwaves, the thermalization of alpha particles in DT plasmas, the study of runaway electrons and ions in tokamaks, and the performance of two-energy component fusion reactors are other examples where the solution of the Fokker-Planck equation is required.

The problem is to solve a nonlinear, time-dependent partial differential equation for the distribution function of each charged species in the plasma, as functions of six phase space variables (three spatial coordinates and three velocity coordinates). Such an equation, even for a single species, exceeds the capability of any present computer, so several simplifying assumptions are required to treat the problem.

With the advent of much more powerful neutral beams, it is now possible to consider neutral-beam-driven tokamak fusion reactors[8]. For such devices, three operating regimes can be considered: (1) the beam-driven thermonuclear reactor, (2) the two-energy component torus (TCT), and (3) the energetic-ion-reactor, e.g., the counterstreaming ion torus (CIT). In order to study reactors in regimes (2) or (3), a non-linear Fokker-Planck model must be used because most of the fusion energy is produced by beam-beam or beam-plasma reactions. Furthermore, when co and counter injection are used, or major

radius compression is employed, a two velocity-space dimensional Fokker-Planck operator is required[9,10].

An example of an important 3-D (r,v,θ) calculation which is beyond the capabilities of the Cray 1 is the modeling of the transport of electron energy out of a tokamak due to the combined effects of a stochastic magnetic field and a radial ambipolar field coupled to Coulomb collisions. This problem is both nonlinear and essentially 3-D. Using an implicit scheme employing a 3-D ICCG matrix inversion package, assuming a mesh of about 120,000 points (a minimum for a physically reasonable 3-D calculation), and a cost of 1.5×10^{-3} seconds per time step per mesh point on the Cray 1, and assuming that a calculation requires 200 time steps, the amount of Cray 1 computer time required is about 10 hours, generally an unacceptable amount of time for a single run. Incidentally, the total of storage required would be about 50% greater than the matrix size or about 3.4×10^6 words. This could be accommodated on the Cray 2.

SUMMARY

In summary, as the fusion program has advanced rapidly in the last few years with the development of more sophisticated theory and experiment, computational requirements for accuracy and realism have increased to the point that Cray 2 capabilities and beyond are required. New features of the machines will allow vectorization of Monte-Carlo, finite element codes, and others which have been scalar until now, a gain of 10 in speed. When programmed to also take advantage of multiprocessing, they will be another factor of 10 faster. This will make revolutionary changes in the importance of such techniques.

It is not possible to define a performance level that represents the ultimate capability for fusion studies. Each successive generation of supercomputers has been exploited to produce more realistic results. Codes to exploit the new hardware capabilities are typically under development before the hardware is actually installed. It is safe to assert that the fusion computing community can effectively use the best performance that the supercomputer manufacturer's are capable of providing for the forseeable future.

ENERGY RESEARCH

INTRODUCTION

During FY84, the Department of Energy (DOE) created the Energy Sciences Advanced Computation activity and established, as its major program, a supercomputer access program. This program was initiated as the result of various panels which had investigated the availability of modern supercomputer resources to the scientific research community within the U.S. and to the DOE research community in particular. It was found that the current availability of modern supercomputer resources within the U.S. fell far short of the amount of these resources needed by the research community and it was also found that modern supercomputers themselves do not have sufficient capability to address many of the computational needs of this community. During FY84 a requirement analysis was conducted throughout the research community which is funded by the Office of Energy Research (ER), and this analysis verified that several Class VI computer systems would be needed to begin satisfying this supressed demand[11].

The disciplines with supercomputing needs include High Energy Physics, Nuclear Physics, Chemical and Materials Sciences, Engineering and Applied Mathemataical Sciences, Geological and Meteorological Sciences and the Biological and related sciences. Extensive computing requirements in these fields have been identified, however, new problem areas are continually being uncovered and the magnitude of the latest demand for supercomputing in the ER programs is just beginning to be understood.

The purpose of the Energy Sciences Advanced Computation Supercomputer Access Program is to provide nationwide high-speed network access to modern centralized facilities within the constraints of budgetary resources. In order to begin addressing this access problem as quickly and as economically as possible, ER decided to utilize the existing National Magnetic Fusion Energy Computer Center (NMFECC) and its installed high-speed satellite network, described earlier, across all ER programs. Because the NMFECC satellite network was already accessible at many DOE laboratories and universities and because this network provides gateways to other networks, such as ARPANET and

TYMNET, many researchers were able to gain access to the NMFECC facilities with very little lead time and minimal additional cost.

For fiscal year 1985, the Office of Energy Research is funding the Cray X-MP/2 computer system installed at the NMFECC in November 1984, to further expand the availability of supercomputer resources to the non-fusion ER programs. This system addresses the near term capability and capacity needs. The Office of Energy Research is requesting funds to replace this Cray X-MP/2 system with a more advanced Class VII system in FY87 in order to provide the capabilities required[11]. The Class VII system will be acquired through a competitive procurement at a time when U.S. vendors are expected to market at least three systems of this capability.

THE NEED FOR MORE POWERFUL COMPUTERS

Historically, scientists who use supercomputers have constrained their numerical simulations to an average execution time of about ten hours. This constraint reflects the scientist's need to make daily progress. Thus, the amount of complexity incorporated in models is scaled to the computer's ability to produce results in about a ten-hour execution time. The capability of a supercomputer dictates the amount of complexity that can be treated. Because of this limitataion, scientists engaged in large-scale numerical simulation have continually sough bigger and faster computers. Today, scientists engaged in energy research need supercomputers that are up to 200 times faster than state-of-the-art equipment.

In order to understand the requirements for more powerful computers, we must explore the generic reasons for having increased computational speed and storage.

Dimensionality. The real world exists in three space dimensions plus time. If computational models reflected the real world exactly and completely, they would treat all four of these dimensions and other parameters that are equivalent to additional dimensions. With current computers, it is possible to treat two space dimensions and time for some problem types, three space dimensions for others, and three space dimensions plus time for a very limited

set of problems. Speed increases of about a factor of 200 in this decade are needed to allow researchers to solve urgent multidimensional problems that are now intractable.

Resolution. Every region of space contains infinitely many points. Thus, the first step in modeling any natural phenomenon is to approximate the space with a finite set of zones, each of which requires a number of calculations. Increasing the number of zones means we can determine more completely and accurately what is happening in any environment, but the computational time grows very rapidly. For example, in a two-dimensional time-dependent model, the running time grows in proportional to the third power of the increase in resolution; increasing the number of zones by just a factor of 2 would increase the time to complete the problem by a factor of 8. Many complex problems now run up to 100 hours, so it is clear that resolution increases of even relatively small factors can overwhelm the capabilities of current supercomputers.

Physics. All computational models dealing with the frontiers of science and technology make simplfying assumptions about the laws of physics in order to keep the calculations from running too long. In some models, including just one additional physical effect can increase running time by a factor of 10. Faster supercomputers with much larger memories will permit researchers to solve problems that cannot now be economically solved.

Combination effects. Although dimensionality, resolution, and physics each have powerful effects on running time by themselves, the overall needs are derived from combinations of these effects. The highly complex problems now being studied in energy research programs, require computational models with higher dimensionality, and with higher resolution, and with more physics.

HIGH ENERGY AND NUCLEAR PHYSICS

The requirement for computers capable of meeting the data reduction needs of a high energy physics laboratory has, historically, been so great that all other computing requirements could be met without significantly impacting the large central facility. However, in the decade of the '80's, several new

computational needs have appeared which require the unique capabilities of supercomputers and clearly require capabilities presently associated with Class VII systems.

The theoretical high energy physics community represents an important class of users with very large computational needs. This is primarily due to the rapid rise of computational quantum field theory, particularly in numerical studies of lattice gauge theory. To put this development in perspective, it should be noted that computer simulation is a generic numerical tool for studying the behavior of particles and fields, and its importance does not rest on any particular fashion nor on the currency of any particular theoretical idea. The ability to carry out such calculations is primarily a result of the rapid increase in available computer power, and as such, it represents a permanent change in the way theoretical physics is done. The needs here fall into two distinct categories. The first category includes the more traditional forms of theoretical computation such as numerical integration, solution of integral or differential equations, calculation of Feynman diagrams, etc. The second category is the large scale numerical simulation of quantum field theory on a lattice. These calculations are highly CPU intensive. The lattice gauge theory algorithms are relatively simple, repetitive and easily vectorizable. Thus, they are well suited to a variety of parallel and pipelined architectures provided that a large, faster accessed memory is available. Even low statistics calculations on modest sized lattices require the equivalent of tens of CRAY hours.

Two newly emerging needs for computer power beyond the scope of Class VI systems are from the accelerator and experiment design communities. An example of an accelerator design requirement is for the turn-by-turn simulation of potential designs for the new superconducting super collider (SSC) accelerator currently in conceptual design. The integrated time needs here are CPU times measured in CRAY-1 equivalent years.

An example of the experiment-design-related requirement is the full simulation of Monte Carlo events in a colliding beam detector system. The number of simulated events run should, ideally, be substantially greater than the number of real physics events to be analyzed. Furthermore, since

experimental results may change the way a detector is tuned, it may be necessary to make the simulations concurrently with the taking of data, i.e., when the data reduction computers are most fully loaded.

Experimental high energy physics data reduction, which has heretofore used standard general purpose computers, also needs a new genertion of computers. The generation of detectors now just coming into use necessarily gather data at very high rates in order to extract the physics of interest from the enormously large accompanying backgrounds. The volume of data collected from these new experiments is several orders of magnitude larger than in experiments performed in the 1980 period. The computational problems are enormous and new classes of supercomputers along with special purpose processors appear to be the only practical way in which to satisfy these unfilled computational needs.

Monte Carlo simulations of lattice gauge theories, and more specifically of Quantum Chromo-dynamics (QCD), while not being the only calculations of interest of particle theory, are presently the most demanding in computational resources and the most likely to produce quantitative predictions. Two essential elements enter into these calculations:

a) the generation of gauge field configurations distributed according to the exp{-S} measure;
b) the evaluation of quark propagators in the background of the gauge fields provided by the above configurations.

The degrees of freedom are made discrete by introduction of a (usually) hypercubical lattice. A lattice extending for n_s sites in the spatial directions and n_t sites in the temporal one entails $4n_s^3 n_t$ gauge dynamical variables associated with the links of the lattice. At present, we lack a quantitative understanding of any collective excitation which may dominate the functional integrals. Therefore all of the above link variables must be treated on the same footing. The lattice must be sufficiently big to contain a hadron, and provide enough resolution so that a lattice with the same physical volume but a finer subdivision would lead to essentially unmodified results (notion of scaling toward the continuum limit). Let us assume, to fix ideas, that the lattice extends for 10 sites in all spatial directions and 20 in the

temporal one. This give a total of 80,000 link variables, i.e., 80,000 SU(3) matrices which must be kept in the memory of the computer for a simulation of QCD. Thus a lattice configuration corresponds to 80,000 x 18 real numbers = 1,440,000 words of memory. To proceed from one configuration to the next all 80,000 link gauge variables must be "upgraded". The upgrading of a single variable involves on the order of 4000 arithmetic operations. We thus obtain an operation count of $\approx$ 320 million to generate a new configuration. Typically hundreds or thousands of configurations must be generated to produce meaningful results. The calculations of the quark propagators are about as demanding in computational resources.

In conclusion a computing center which would serve the interest of high energy theorists should be endowed with one of the most powerful mainframes available, both in computational speed and in memory size, such as the ER Class VII system proposed for FY87.

BASIC ENERGY SCIENCE

Material Sciences

The development and proliferation of investigations of diverse material systems and phenomena via computer simulation and modeling is a rich field of scientific endeavour anchored in the physical sciences (with cross-fertilization links to advances in applied mathematics and computer science), made possible singularly by the advent of high-powered computers. Computer simulations provide information about phenomena and processes in material systems with refined microscopic spatial and temporal resolution and enable investigations of the dynamical evolution of complex systems under extreme conditions where data from experiments or other methods of investigation is not attainable. In addition such studies provide benchmarks for critical testing and refinement of theoretical concepts and aid in the interpretation of experimental observations.

Current simulation methods involve the generation and analysis of phase-space trajectories of an interacting many-particle system either by the direct numerical integration of the equations of motion (molecular dynamics-MD,

and reaction-trajectory-TJ, methods) or via the sampling of phase-space configurations (Monte Carlo-MC). In either case the many-body nature of the systems under study and the statistical modes of analyses dictate the necessity for extended computer time and storage capabilities.

The wide range of materials system investigated by computer simulations include: the equilibrim and non-equilibrium structure and dynamics of materials at different states of aggregation (solids and liquids) and the kinetics and dynamics of phase transformations; properties of metastable systems (supercooled liquids, quenched liquids, gasses); homo and multicomponent materials; ordered versus disordered (amorphous) solids; surfaces; interfaces and inter-phase interfaces, i.e, solid-solid (superlattices and coherent structures), solid-liquid (epitaxial crystal growth and homogeneous nucleation), solid-gas (molecular beam epitaxy, heterogeneous catalysis).

Simulation studies on these systems allow investigation of structural and dynamical characteristics, kinetics and dynamics of phase-transformations, transport and non-linear phenomena (heat, matter, electrical), diffusion processes and reaction dynamics. Furthermore modifications of the intrinsic properties of condensed matter systems and phenomena (such as fracture, solid transformations, plastic flow), due to external fields (mechanical stress, heat gradient etc.) can be investigated. In addition to an improved understanding of existing material systems, simulation studies could serve as the impetus for exploration of methods of preparation and growth of novel materials.

Underlying simulation studies of extended condensed matter systems is the notion that the properties of the "calculational sample" on which the simulation is carried out, extended via the commonly used periodic boundary conditions, are a faithful representation of the nature of the macroscopic system. Among the factors which dictate the size of the "calculational sample" are the ranges of interparticle interaction potentials and fluctuation wavelengths. Thus for example the MD simulation of the structural and dynamical properties of a solid simple metal (e.g., Al) requires a system containing ~2000 particles; the simulation of binary liquid metals and supercooled liquids require an even larger sample due to concentration fluctuations. Simulations of stressed crystals, fracture and plastic flow,

shock wave propagation, the dynamics of melting and hydrodynamical phenomena would require systems where the number of particles would be 5000-10,000. It should be noted that in the presence of long range and realistic multibody forces the computing time grows as a (low) power of the number of particles. Such requirements necessitate memory capacity beyond CRAY-1S capability and large increases in computational speed.

A critical input in materials simulations is the interparticle interaction potential. A faithful simulation requires the calculation of such potentials via pseudo-potential methods which, for metalic systems, depends upon the thermodynamic state variables (density, temperature, pressure). Simulations of nonequilibrium phenomena (such as solidification, quenching etc.) in which the state variables themselves evolve in time require a self-consistent adjustment of the interaction potentials along with the dynamical evolution of the system.

The coupled complexities of size and interaction potential calculations make such simulations prohibitive on the Cray 1. Furthermore, the magnitude of such simulations dictate substantial time requirements, for example, 80 minutes of CRAY-1 time allow the generation of ~5000 integration time steps for a system containing 1500 particles interacting via simple truncated Lennard-Jones potentials, with a fully optimized code. Note that this is the least demanding model from a computational point of view. A typical study of the solidification of such a system requires 50,000 integration time steps. It should be emphasized that the above considerations are dictated by the nature of the physical systems and phenomena and cannot be compromised by approximate treatments which will prejudice and distort the simulation results. Thus progress in this field can be made with substantial access to the Class VII computing facilities proposed for FY87.

REFERENCES

1. B.I. Cohen and J. Killeen, Phys. Today (May 1983).

2. M.N. Rosenbluth, D.A. Monticello, H.R. Strauss, R.D. White, Phys. Fluids 19, p. 1987 (1976).

3. D. Schnack and J. Killeen, Nucl. Fusion 19, p. 877 (1979).

4. E.J. Caramana, R.A. Nebel, D. Schnack, Phys. Fluids 25, p. 1305 (1983).

5. J. Dawson, H. Okuda, and B. Rosen, Methods in Comp. Phys. 16, p. 282 (1976).

6. A.B. Langdon and B. Lasinski, Methods in Comp. Phys. 16, p. 327 (1976).

7. C. Nielson and H. Lewis, Methods in Comp. Phys. 16, p. 367 (1976).

8. D.L. Jassby, Nucl. Fusion 17, p. 309 (1977).

9. J. Killeen, A.A. Mirin, and M.E. Rensink, Methods in Comp. Phys. 16, p. 389 (1976).

10. J. Killeen, A.A. Mirin, and M.G. McCoy, Modern Plasma Physics, IAEA-SMR-61/110, p. 395 (1981).

11. "The Role of Supercomputers in Energy Research Programs," U.S. Department of Energy, DOE/ER-0218, February 1985.